

Structured Deep Hashing with Convolutional Neural Networks for Fast Person Re-identification

Lin Wu[‡], Yang Wang[†]

[‡]ISSR, ITEE, The University of Queensland, Brisbane, QLD, 4072, Australia

[†]The University of New South Wales, Kensington, Sydney, Australia
lin.wu@uq.edu.au wangi@cse.unsw.edu.au

Abstract

Given a pedestrian image as a query, the purpose of person re-identification is to identify the correct match from a large collection of gallery images depicting the same person captured by disjoint camera views. The critical challenge is how to construct a robust yet discriminative feature representation to capture the compounded variations in pedestrian appearance. To this end, deep learning methods have been proposed to extract hierarchical features against extreme variability of appearance. However, existing methods in this category generally neglect the efficiency in the matching stage whereas the searching speed of a re-identification system is crucial in real-world applications. In this paper, we present a novel deep hashing framework with Convolutional Neural Networks (CNNs) for fast person re-identification. Technically, we simultaneously learn both CNN features and hash functions/codes to get robust yet discriminative features and similarity-preserving hash codes. Thereby, person re-identification can be resolved by efficiently computing and ranking the Hamming distances between images. A structured loss function defined over positive pairs and hard negatives is proposed to formulate a novel optimization problem so that fast convergence and more stable optimized solution can be obtained. Extensive experiments on two benchmarks CUHK03 [Li *et al.*, 2014] and Market-1501 [Zheng *et al.*, 2015] show that the proposed deep architecture is efficacy over state-of-the-arts.

1 Introduction

re-identification is task of matching persons observed from non-overlapping camera views based on visual appearance. It has gained considerable popularity in video surveillance, multimedia, and security system by its prospect of searching a person of interest from a large amount of video sequences [Wang *et al.*, 2016c; Ye *et al.*, 2016; Sunderrajan and Manjunath, 2016]. The major challenge arises from the variations in

human appearances, poses, viewpoints and background cluster across camera views. Some examples are shown in Fig.???. Towards this end, many approaches [Farenzena *et al.*, 2010; Yang *et al.*, 2014; Zhao *et al.*, 2014; Pedagadi *et al.*, 2013; Li *et al.*, 2013; Paisitkriangkrai *et al.*, 2015] have been proposed by developing a combination of low-level features (including color histogram [Gray and Tao, 2008], spatial co-occurrence representation [Wang *et al.*, 2007], LBP [Xiong *et al.*, 2014] and color SIFT [Zhao *et al.*, 2013]) against variations (*e.g.*, poses and illumination) in pedestrian images. However, these hand-crafted features are still not discriminative and reliable to such severe variations and misalignment across camera views.

Recently, deep learning methods [Li *et al.*, 2014; Ahmed *et al.*, 2015; Yi *et al.*, 2014; Ding *et al.*, 2015; Wu *et al.*, 2016b; Chen *et al.*, 2016; Xiao *et al.*, 2016] have been proposed to address the problem of person re-identification by learning deeply discriminative Convolutional Neural Network (CNN) features in a *feed-forward* and *back-propagation* manner. It extracts hierarchical CNN features from pedestrian images; the subsequent metric-cost part compares the CNN features with a chosen metric encoded by specific loss functions, *e.g.*, contrastive (pair-wise) [Li *et al.*, 2014; Ahmed *et al.*, 2015; Wu *et al.*, 2016b] or triplet [Yi *et al.*, 2014; Chen *et al.*, 2016] loss functions. However, such typical deep learning methods are not efficient in real-time scenario, due to the less-efficiency of matching two pedestrian images by extracting and comparing hierarchical CNN features. In fact, the excellent recognition accuracy in neural network-based architectures comes at expense of high computational cost both at training and testing time. The main computational expense for these deep models comes from convolving filter maps with the entire input image, making their computational complexity at least linear in the number of pixels. And matching these CNN features to obtain similarity values is not fast enough to be applicable in real-world applications. In this paper, we aim to reduce the computational burden of person re-identification by developing a fast re-identification framework.

1.1 Motivation

To cope with ever-growing amounts of visual data, deep learning based hashing methods have been proposed to simultaneously learn similarity-preserved hashing functions and discriminative image representation via a deep architecture [Lai *et al.*, 2015; Zhao *et al.*, 2015; Zhang *et al.*, 2015]. Simply delving existing deep hashing approaches into a person re-identification system is not trivial due to the difficulty of generalizing these pre-trained models to match pedestrian images in disjoint views. Fine-tuning is a plausible way to make pre-trained models suitable to re-identification, however, to suit their models, training images are commonly divided into mini-batches, where each mini-batch contains a set of *randomly* sampled positive/negative pairs or triplets. Thus, a contrastive or triplet loss is computed from each mini-batch, and the networks try to minimize the loss function feed-forwardly and update the parameters through back-propagation by using Stochastic Gradient Decent (SGD) [Wilson and Martinez, 2003].

We remark that randomly sampled pairs/triplets carry little helpful information to SGD. For instance, many triplet units can easily satisfy the relative comparison constraint in a triplet loss function (Eq (3)), resulting into a slow convergence rate in the training stage. Worse still, mini-batches with random samples may fail to obtain a stable solution or collapsed into a local optimum if a contrastive/triplet loss function is optimized [Song *et al.*, 2016]. To this end, a suitable loss function is highly demanded to work well with SGD over mini-batches.

In this paper, we propose a deep hashing based on CNNs to efficiently address the problem of person re-identification. To mitigate the undesirable effects caused by contrastive/triplet loss function, we propose a structured loss function by *actively* adding hard negative samples into mini-batches, leading to a structured deep hashing framework. The proposed structured loss can guide sub-gradient computing in SGD to have correct directions, and thus achieves a fast convergence in training.

1.2 Our Approach

One may easily generate a straightforward two-stage deep hashing strategy by firstly extracting CNN features from a pre-trained model *e.g.*, AlexNet [Krizhevsky *et al.*, 2012], followed by performing the learned hash functions (separate projection and quantization step) to convert such CNN features into binary codes. However, as demonstrated in section 4, such a strategy cannot obtain optimal binary codes. As such binary codes may not well characterize the supervised information from training data *i.e.*, intra-personal variation and inter-personal difference, due to the *independence* of two stages. In fact, such two stages can boost each other to achieve much better performance, that is, the learned binary codes can guide the learning of useful CNN features, while CNN features can in turn help learn semantically similarity-preserving hash function/codes.

Motivated by this, we present a structured deep hashing architecture to *jointly* learn feature representations and hash codes for person re-identification. The overall framework is illustrated in Fig.1. In our architecture, mini-batches contain




Figure 1: Overview of our deep hashing framework for person re-identification. Our deep neural network takes a feed-forward, back-propagation strategy to learn features and hash codes simultaneously. During the feed-forward stage, the proposed network performs inference from a mini-batch. The mini-batch is put through a stack of convolutional layers to generate nonlinear yet discriminative features, which are subsequently mapped to output feature vector by fully-connected layers (FC). Meanwhile, a hash function layer is introduced at the top of FC layer to learn hash codes that are optimized by a structured loss function to preserve their similarities/dissimilarities. In back-propagation, parameters are updated by computing their Stochastic Gradient Decent (SGD) w.r.t. the mini-batch.

all positive pairs for a particular pedestrian, meanwhile each positive pair (has a query image and its correct match image from a different camera view) is augmented by actively selected hard negatives for its query and match image, respectively. Such mini-batches are taken into the inputs of deep network with a structured loss function optimized to learn CNN features and hash functions jointly.

The major contributions are summarized below:

- To the best of our knowledge, we are the first to solve person re-identification efficiently by presenting a structured deep hashing model. This makes our paper distinct from existing studies [Wang *et al.*, 2016c; Ye *et al.*, 2016; Sunderrajan and Manjunath, 2016] where the matching efficiency is not addressed.
- By simultaneously learning CNN features and hash functions/codes, we are able to get robust yet discriminative features against complex pedestrian appearance and boosted hash codes, so that every two deep hashing codes learned from the same identities are close to each other while those from different identities are kept away.
- To combat the drawbacks of the contrastive/triplet loss, we propose a structured loss function where mini-batches are augmented by considering hard negatives. Also, the proposed structured loss function that is imposed at the top layer of the network can achieve fast convergence and a stable optimized solution.

2 Related Work

In this section, we briefly review deep learning based on CNNs for person re-identification and several typical hashing methods, as they are closely related to our proposed technique.

In literature of person re-identification, many studies try to address this challenging problem by either seeking a robust feature representation [Farenzena *et al.*, 2010; Zhao *et al.*, 2014; Wang *et al.*, 2013; Wu *et al.*, 2013; Wang *et al.*, 2016a; Wang *et al.*, 2015a; Wang *et al.*, 2015c; Wang *et al.*, 2017b; Wang *et al.*, 2016b; Gray and Tao, 2008; Zhao *et al.*, 2013] or casting it as a metric learning problem where more discriminative distance metrics are learned to handle features extracted from person images across camera views [Li *et al.*, 2013; Kostinger *et al.*, 2012; Pedagadi *et al.*, 2013; Xiong *et al.*, 2014; Liao *et al.*, 2015; Zhang *et al.*, 2016]. The first aspect considers to find features that are robust to challenging factors while preserving identity information. The second stream generally tries to minimize the intra-class distance while maximize the inter-class distance. Also, person re-identification can be approached by a pipeline of image search where a Bag-of-words [Zheng *et al.*, 2015] model is constructed to represent each pedestrian image and visual matching refinement strategies can be applied to improve the matching precision. Readers are kindly referred to [Gong *et al.*, 2014] to have more reviews.

A notable improvement on person re-identification is achieved by using Convolutional Neural Networks (CNNs) [Li *et al.*, 2014; Ahmed *et al.*, 2015; Yi *et al.*, 2014; Ding *et al.*, 2015; Wu *et al.*, 2016b; Wu *et al.*, 2016a; Chen *et al.*, 2016; Wang *et al.*, 2017a; Xiao *et al.*, 2016], which can jointly learn robust yet discriminative feature representation and its corresponding similarity value in an end-to-end fashion. However, existing deep learning methods in person re-identification are facing a major challenge of efficiency, where computational time required to process an input image is very high due to the convolution operations with the entire input through deep nets. Thus, from a pragmatical perspective, an advanced yet fast neural network-based architecture is highly demanded. This motivated us to develop an efficient deep learning model to alleviate the computational burden in person re-identification.

Hashing is an efficient technology in approximate nearest neighbor search with low storage cost of loading hash codes. Learning-based hash methods can be roughly divided into two categories: unsupervised methods and supervised methods. Unsupervised methods including Spectral Hashing [Weiss *et al.*, 2008; Wang *et al.*, 2015b] and Iterative Quantization [Gong and Lazebnik, 2011] only use the training data to learn hash functions. Supervised methods try to leverage supervised information to learn compact binary codes. Some representative methods are Binary Reconstruction Embedding (BRE) [Kulis and Darrell, 2009], Minimal Loss Hashing (MLH) [Norouzi and Blei, 2011], and Supervised Hashing with Kernels (KSH) [Liu *et al.*, 2012].

Nonetheless, these hashing methods often cope with images represented by hand-crafted features (*e.g.*, SIFT [Zhao *et al.*, 2013]), which are extracted before projection and quanti-




Figure 2: Illustrations on different loss functions. (a) Contrastive loss; (b) Triplet ranking loss; (c) Our structured loss. Here, x 's and y 's indicate hash codes of pedestrian images captured by query and gallery camera view, respectively. For a specific pedestrian's hash codes x_i , its correct match's code is y_i from a different view. Green edges and red edges represent similar and dissimilar examples, respectively. Our method explicitly adds hard negatives (*e.g.*, y_m, y_n) for all positive pairs (*e.g.*, (x_1, y_1) and (x_2, y_2)) into mini-batches.

zation steps. Moreover, they usually seek a linear projection which cannot capture the nonlinear relationship of pedestrian image samples¹. Even though some kernel-based hashing approaches [Liu *et al.*, 2012; Wu and Wang, 2017] have been proposed, they are stricken with the efficiency issue. To capture the non-linear relationship between data samples while keeping efficient, Liong *et al.* [Liong *et al.*, 2015] present a Deep Hashing to learn multiple hierarchical nonlinear transformation which maps original images to compact binary code and thus supports large-scale image retrieval. A supervised version named Semantic Deep Hashing is also presented in [Liong *et al.*, 2015] where a discriminative item is introduced into the objective function. However, the above methods did not include a pre-training stage in their networks, which may make the generated hash codes less semantic. To keep the hash codes semantic, Xia *et al.* [Xia *et al.*, 2014] proposed a deep hashing architecture based on CNNs, where the learning process is decomposed into a stage of learning approximate hash codes from supervised priors, which are used to guide a stage of simultaneously learning hash functions and image representations.

More recently, to generate the binary hash codes directly from raw images, deep CNNs are utilized to train the model in an end-to-end manner where discriminative features and hash functions are simultaneously optimized [Lai *et al.*, 2015; Zhao *et al.*, 2015; Zhang *et al.*, 2015]. However, in training stage, they commonly take mini-batches with randomly sampled triplets as inputs, which may lead to local optimum or unstable optimized solution.

By contrast, in this paper we deliver the first efforts in proposing a structured deep hashing model for person re-identification, which allows us to jointly learn deep feature representations and binary codes faithfully. The proposed structured loss function benefits us from achieving fast convergence and more stable optimized solutions, compared with pairwise/triplet ranking loss.

¹Pedestrian images typically undergo compounded variations in the form of human appearance, view angles, and human poses.

3 Structured Deep Hashing for Person Re-identification

Our major contribution is to jointly learn feature representation from raw person images and their mappings to hash codes by presenting an improved deep neural network. The proposed network takes a mini-batch as its input which contains images in a form of positive/negative pairs. The architecture consists of three components: 1) a stack of convolution layers followed by max pooling to learn non-linear feature mappings from raw pedestrian images; 2) a hash layer connected to the first and the second fully connected layers; 3) a structured loss function is designed to optimize the whole mini-batch. The architecture overview is illustrated in Fig.1.

3.1 Learning Deep Hashing Functions

Assuming \mathcal{I} to be the original image space, a hash function $f: \mathcal{I} \rightarrow \{0, 1\}^r$ is treated as a mapping that projects an input image \mathcal{I} into a r -bit binary code $f(\mathcal{I})$ while preserving the similarities of person images across camera views.

Learning based hashing methods aim to seek a set of hash functions to map and quantize each sample into a compact binary code vector. Assuming we have r hash functions to be learned, which map an image \mathcal{I} into a r -bit binary code vector $f(\mathcal{I}) = [f_1(\mathcal{I}), f_2(\mathcal{I}), \dots, f_r(\mathcal{I})]$. Although many learning-based hashing methods have been proposed [Gong *et al.*, 2012; Gong and Lazebnik, 2011; He *et al.*, 2013; Norouzi and Blei, 2011; Kulis and Darrell, 2009], most of them essentially learn a single linear projection matrix, which can not well capture the nonlinear relationship of samples. Admittedly, some kernel-based hashing methods are available [Liu *et al.*, 2012; He *et al.*, 2010], they instead suffer from the efficiency issue because kernel-based methods cannot have explicit nonlinear mapping.

In this work, we propose to learn deep hash functions with CNNs to jointly learn feature representation from raw pixels of pedestrian images and their mappings to hash codes. In this way, feature representations for person images can be learned more optimally compatible with the coding process, thus producing optimal hash codes.

During training, the input to our network is a mini-batch containing pairs of fixed-size 160×60 RGB images. The images are passed through four convolutional layers, where we use filters with a very small receptive field: 3×3. The convolution stride is fixed to 1 pixel. Spatial pooling is carried out by three max-pooling layers. Max-pooling is performed over a 2×2 pixel window, with stride 2. After a stack of convolution layers, we have two fully-connected layers where the first one has 4096 dimension and the second is 512-dim, which are then fed into the hash layer to generate a compact binary code. We show details of layers in CNNs in Table 1.

Inspired by [Sun *et al.*, 2014], we add a bypass connection between the first fully connected layer and the hash layer to reduce the possible information loss. Another reason is features from the second fully connected layer is very semantic and invariant, which is unable to capture the subtle difference between person images. Thus, we rewrite the deep hash function as:

$$f(\mathcal{I}, \mathbf{w}_i) = \text{sigmoid}(\mathbf{w}_i^T [g_1(\mathcal{I}); g_2(\mathcal{I})]), \quad (1)$$

Table 1: Layer parameters of convolutional neural networks. The output dimension is given by height×width×width. FS: filter size for convolutions. Layer types: C: convolution, MP: max-pooling, FC: fully-connected. All convolution and FC layers use hyperbolic tangent as activation function.

Name	Type	Output Dim	FS	Sride
Conv0	C	157 × 57 × 32	3×3	1
Pool0	MP	79 × 29 × 32	2×2	2
Conv1	C	76 × 26 × 32	3×3	1
Pool1	MP	38 × 13 × 32	2×2	2
Conv2	C	35 × 10 × 32	3×3	1
Pool2	C	18 × 5 × 32	3×3	2
Conv3	C	15 × 2 × 32	3×3	1
Pool4	MP	15 × 2 × 32	1×1	1
FC1	FC	-	4096	-
FC2	FC	-	512	-

where $\text{sigmoid}(t) = 1/(1 + \exp(-w^T t))$, \mathbf{w}_i denotes the weights in the i -th hash function, $g_1(\cdot)$ and $g_2(\cdot)$ represent feature vectors from the outputs of the two fully connected layers, respectively. Then, we have $f(\mathcal{I}, \mathbf{W}) = [f(\mathcal{I}, \mathbf{w}_1), \dots, f(\mathcal{I}, \mathbf{w}_r)]$. After the deep architecture is trained, the hashing code for a new image \mathcal{I} can be done by a simple quantization $b = \text{sign}(f(\mathcal{I}, \mathbf{W}))$, where $\text{sign}(v)$ is a sign function on vectors that for $i = 1, 2, \dots, r$, $\text{sign}(v_i) = 1$ if $v_i > 0$, otherwise $\text{sign}(v_i) = 0$.

3.2 Structured Loss Optimization

In deep metric learning for person re-identification, the network is often trained on data in the form of pairs [Li *et al.*, 2014; Ahmed *et al.*, 2015; Yi *et al.*, 2014] or triplet ranking [Ding *et al.*, 2015]. Thus, there are two commonly used cost functions, contrastive/pairwise loss and triplet ranking loss, which can be used in hash code optimization. We briefly revisit the two loss functions and then introduce the proposed structured loss function.

Contrastive/Pairwise Loss Function Given a person’s binary codes \mathbf{x}_i and its correct match’s codes \mathbf{y}_i from a different camera view, the contrastive training tries to minimize the Hamming distance between a positive pair of $(\mathbf{x}_i, \mathbf{y}_i)$ and penalize the negative pairs $(\mathbf{x}_i, \mathbf{y}_j)$ ($i \neq j$) with a Hamming distance smaller than a margin. The contrastive cost function can be defined as

$$F = \sum_{(i,j)} a_{i,j} \|\mathbf{x}_i - \mathbf{y}_j\|_H + (1 - a_{i,j}) (\max(0, 1 - \|\mathbf{x}_i - \mathbf{y}_j\|_H)) \quad (2)$$

where $\mathbf{x}_i, \mathbf{y}_j \in \{0, 1\}^r$ and $\|\cdot\|_H$ represents the Hamming distance. The label $a_{i,j} \in \{0, 1\}$ indicates whether a pair of binary codes $(\mathbf{x}_i, \mathbf{y}_j)$ depicting the same person.

Triplet Ranking Loss Function Some recent studies have been made to learn hash functions that preserve relative similarities of the form of a triplet data $(\mathcal{I}, \mathcal{I}^+, \mathcal{I}^-)$ where image \mathcal{I} (anchor) of a specific person is more similar to all other

images \mathcal{I}^+ (positive) of the same person than it is to any image \mathcal{I}^- (negative) of any other person (images \mathcal{I}^+ and \mathcal{I}^- are from a different camera view from \mathcal{I}). Specifically, in hash function learning, the goal is to find a mapping $f(\cdot)$ such that the binary code $f(\mathcal{I}) = \mathbf{x}_i$ is closer to $f(\mathcal{I}^+) = \mathbf{y}_i$ than to $f(\mathcal{I}^-) = \mathbf{y}_j$ ($j \neq i$). Thus, we want

$$\|\mathbf{x}_i - \mathbf{y}_i\|_H + 1 < \|\mathbf{x}_i - \mathbf{y}_j\|_H, \forall (\mathbf{x}_i, \mathbf{y}_i, \mathbf{y}_j) \in \mathbb{T}, \quad (3)$$

where \mathbb{T} is the set of all possible triplets in the training set and has cardinality N . Accordingly, the triplet ranking hinge loss is defined by

$$F = \sum_i \max(0, 1 - (\|\mathbf{x}_i - \mathbf{y}_j\|_H - \|\mathbf{x}_i - \mathbf{y}_i\|_H)) \quad (4)$$

s.t. $\mathbf{x}_i, \mathbf{y}_i, \mathbf{y}_j \in \{0, 1\}^r$.

A noticeable difference between a contrastive embedding and a triplet embedding is that a triplet unit with similar and dissimilar inputs provide some distance comparison context for the optimization process, as opposed to the contrastive loss that the network minimizes (same class) or maximizes (different classes) as much as possible for each pair independently [Hoffer and Ailon, 2014].

In triplet embedding, however, generating all possible triplets would result in many triplets that easily fulfill the constraint in Eq (3), which is known as over-sampling. These triplets would not contribute to the training whereas resulting in slow convergence. An alternative strategy is to perform a smart sampling where one must be careful to avoid too much focus on hard training exemplars due to the possibility of over-fitting. Thus, it is crucial to *actively* select informative hard exemplars in order to improve the model.

Below, we introduce our structured loss which can avoid aforementioned over or under-sampling dilemmas by virtue of actively adding difficult neighbors to positive pairs into training batches.

The Proposed Structured Loss Function Previous works on person re-identification implement a Stochastic Gradient Decent (SGD) [Wilson and Martinez, 2003] by drawing pairs or triplets of images uniformly at random. They didn't fully makes use of the information of the mini-batch that is sampled at a time and not only individual pairs or triplets. By contrast, we propose a structured loss over a mini-batch in order to take fully advantage of the training batches used in SGD. Meanwhile, the proposed structured loss can ensure fast convergence and stableness in training.

As shown in Fig.2 (c), the structured loss is conducted on all positive pairs and corresponding close ("difficult") negative pairs across camera views. Specifically, it can be formulated as

$$F = \frac{1}{|\bar{P}|} \sum_{\mathbf{x}_i, \mathbf{y}_i \in \bar{P}} \max(0, F_{\mathbf{x}_i, \mathbf{y}_i}),$$

$$F_{\mathbf{x}_i, \mathbf{y}_i} = \max(\max(0, 1 - \|\mathbf{x}_i - \mathbf{y}_k\|_H), \max(0, 1 - \|\mathbf{y}_i - \mathbf{y}_l\|_H) + \|\mathbf{x}_i - \mathbf{y}_i\|_H),$$

s.t. $\mathbf{x}_i, \mathbf{y}_i, \mathbf{y}_k, \mathbf{y}_l \in \{0, 1\}^r, (\mathbf{x}_i, \mathbf{y}_k) \in \bar{N}, (\mathbf{y}_i, \mathbf{y}_l) \in \bar{N},$ (5)

where \bar{P} and \bar{N} denote the set of positive and negative pairs in each mini-batch. The process of selecting positive and negative samples is elaborated in Section 3.3.

Difference to contrastive and triplet ranking loss:

- In pairwise training with $O(m)$ separate pairs in the batch, a total of $O(m^2)$ pairs can be generated accordingly. However, these negative edges induced between randomly sampled pairs carry very limited information [Song *et al.*, 2016]. By contrast, selected difficult exemplars are sharper cases that a full sub-gradient method would more likely focus on;
- Compared with triplet embedding containing randomly sampled triplets, our training batch is augmented by adding negative neighbors bilaterally for each positive pairs. By doing this, the optimization process is conducted on most violate constraints, leading to fast convergence.

Fig.2 (a) and (b) illustrates a batch of positive/negative pairs and triplets with corresponding contrastive loss and triplet ranking loss. Green edges represent positive pairs (the same person) and red edges represent negative pairs (different individuals). Please note that these pairs and triplets are sampled completely random into a mini-batch. Fig.2 (c) illustrates the mining process for two positive pairs in the batch where for each image in a positive pair we seek its close (hard) negative images. We can see that our method allows mining the hard negatives from both the query image (*e.g.*, \mathbf{x}_1) and its correct match (*e.g.*, \mathbf{y}_1) of a pair against gallery images (*e.g.*, $\mathbf{y}_m, m \neq 1$).

Optimization For ease of optimization, we relax Eq.(5) by replacing the Hamming norm with the ℓ_2 -norm and replacing the integer constraints on \mathbf{x} 's and \mathbf{y} 's with the range constraints. The modified loss function is

$$F = \frac{1}{|\bar{P}|} \sum_{\mathbf{x}_i, \mathbf{y}_i \in \bar{P}} \max(0, F_{\mathbf{x}_i, \mathbf{y}_i}),$$

$$F_{\mathbf{x}_i, \mathbf{y}_i} = \max(\max(0, 1 - \|\mathbf{x}_i - \mathbf{y}_k\|_2^2), \max(0, 1 - \|\mathbf{y}_i - \mathbf{y}_l\|_2^2) + \|\mathbf{x}_i - \mathbf{y}_i\|_2^2),$$

s.t. $\mathbf{x}_i, \mathbf{y}_i, \mathbf{y}_k, \mathbf{y}_l \in [0, 1]^r, (\mathbf{x}_i, \mathbf{y}_k) \in \bar{N}, (\mathbf{y}_i, \mathbf{y}_l) \in \bar{N}.$ (6)

The variant of structured loss is convex. Its sub-gradients with respect to \mathbf{x}_i , \mathbf{y}_i , \mathbf{y}_k , and \mathbf{y}_l are

$$\frac{\partial F}{\partial \mathbf{x}_i} = (2\mathbf{y}_k - 2\mathbf{y}_i) \times \mathbb{I}[2 + \|\mathbf{x}_i - \mathbf{y}_i\|_2^2 > \|\mathbf{x}_i - \mathbf{y}_k\|_2^2 + \|\mathbf{y}_i - \mathbf{y}_l\|_2^2]$$

$$\frac{\partial F}{\partial \mathbf{y}_i} = (2\mathbf{y}_l - 2\mathbf{x}_i) \times \mathbb{I}[2 + \|\mathbf{x}_i - \mathbf{y}_i\|_2^2 > \|\mathbf{x}_i - \mathbf{y}_k\|_2^2 + \|\mathbf{y}_i - \mathbf{y}_l\|_2^2]$$

$$\frac{\partial F}{\partial \mathbf{y}_k} = 2\mathbf{x}_i \times \mathbb{I}[2 + \|\mathbf{x}_i - \mathbf{y}_i\|_2^2 > \|\mathbf{x}_i - \mathbf{y}_k\|_2^2 + \|\mathbf{y}_i - \mathbf{y}_l\|_2^2]$$

$$\frac{\partial F}{\partial \mathbf{y}_l} = 2\mathbf{y}_i \times \mathbb{I}[2 + \|\mathbf{x}_i - \mathbf{y}_i\|_2^2 > \|\mathbf{x}_i - \mathbf{y}_k\|_2^2 + \|\mathbf{y}_i - \mathbf{y}_l\|_2^2]$$

(7)

The indicator function $\mathbb{I}[\cdot]$ is the indicator function which outputs 1 if the expression evaluates to true and outputs 0 otherwise. Thus, the loss function in Eq.(5) can be easily integrated into back propagation of neural networks. We can see that our structured loss provides informative gradients signals for all negative pairs which are within the margin of any positive pairs. In contrast to existing networks like [Li *et al.*, 2014; Ahmed *et al.*, 2015] where only hardest negative gradients are updated, making the training easily over-fit, the proposed structured loss makes the optimization much more stable.

3.3 Hard Negative Mining for Mini-batches

As mentioned before, our approach differs from existing deep methods by making full information of the mini-batch that is sampled at a time, including positive pairs and their difficult neighbors. Please note that difficult neighbors are defined only with respect to the gallery camera view. The motivation of doing this is to enhance the mini-batch optimization in network training because the sub-gradient of F_{x_i, y_i} would use the close negative pairs. Thus, our approach biases the sample towards including “difficult” pairs.

In this paper, we particularly select a few positive pairs at random, and then **actively** add their difficult (hard) neighbors into the training mini-batch. This augmentation adds relevant information that a sub-gradient would use. Specifically, we determine the elements in mini-batches by on-line generation where all anchor-positive pairs in any identity are kept while selecting the hard negatives for both the anchor and its positive correspondence. In fact, this procedure of mining hard negative edges amounts to computing the loss augmented inference in structured prediction setting [Tschantz *et al.*, 2004; Joachims *et al.*, 2009; Song *et al.*, 2016]. Intuitively, the loss from hard negative pairs should be penalized more heavily than a loss involving other pairs. In this end, our structured loss function contains enough negative examples within the margin bound, which can push the positive examples towards the correct direction and thus making the optimization much more stable.

Example 1 Fig.3 shows failure cases in 2D profile with samples from three different classes, visualized by pink circles, green squares, and magenta triangles, respectively. The contrastive embedding has failure conditioned that randomly sampled negative y_j is collinear² with examples from a third class (purple triangles). For triplet embedding, the degenerated case happens when a negative y_j is within the margin bound with respect to the anchor x_i and its positive y_i . In this situation, both contrastive and triplet embedding incorrectly enforce the gradient direction of positives towards examples from the third class. By contrast, through explicitly mining hard negatives within the margin w.r.t. the positive x_i , the proposed structured embedding can push the positives towards the correct direction.

Theorem 1 *Margin maximization. Hard negative mining on mini-batches is equivalent to computing the loss aug-*

²Three or more points are said to be collinear if they lie on a single straight line.




Figure 3: Illustration on degeneration of contrastive and triplet ranking embedding with randomly sampled training pairs. Pink circles, green squares, and purple triangles indicate three different classes. Dotted blue circles regulate the margin bound where the loss becomes zero out of the bound. Magenta arrows denote the negative sub-gradient direction for positive samples. Left: Contrastive embedding. Middle: Triplet embedding. Right: Proposed structured embedding.

mented inference, which promotes margin maximization in pairwise/triplet units.

Proof. Following the definitions in Eq.(5), the condition of zero training error can be compactly written as a set of non-linear constraints

$$\forall i : \max_{y \in \mathcal{Y} \setminus y_i} \{ \langle w, H(x_i, y) \rangle \} < \langle w, H(x_i, y_i) \rangle. \quad (8)$$

where \mathcal{Y} contains training samples from cross-camera view against x_i . $H(\cdot)$ denotes Hamming distance. Each non-linear inequality in Eq.(8) can be equivalently replaced by $|\mathcal{Y}| - 1$ linear inequalities, and thus we have

$$\forall i, \forall y \in \mathcal{Y} \setminus y_i : \langle w, \delta H_i(y) \rangle < 0; \quad (9)$$

$$\delta H_i(y) \equiv H(x_i, y) - H(x_i, y_i).$$

Recall Eq.(5) that the hard negative mining is equivalent to augmenting the loss as $\bar{H}_i(y) = H(x_i, y) - H(x_i, y_i) + H(y_i, y)$. Thus, the linear constraint in Eq.(9) is updated as

$$\forall i, \forall y \in \mathcal{Y} \setminus y_i : \langle w, \delta \bar{H}_i(y) \rangle < 0; \quad (10)$$

$$\Leftrightarrow \langle w, \delta H_i(y) \rangle + \langle w, H(y_i, y) \rangle < 0.$$

In Eq.(10), since the term $\langle w, H(y_i, y) \rangle \geq 1 - \epsilon_i$, $\epsilon_i \geq 0$ is a small slack variable, the term $\langle w, \delta H_i(y) \rangle$ is imposed a more tight constraint on its margin maximization. \square

4 Experiments

In this section, we conduct extensive evaluations of the proposed architecture on two largest datasets in person re-identification: **CUHK03** [Li *et al.*, 2014] and **Market-1501** [Zheng *et al.*, 2015].

4.1 Experimental Settings

Datasets Person re-identification comes with a number of benchmark datasets such as VIPeR [Gray *et al.*, 2007], PRID2011 [Hirzer *et al.*, 2011], and iLIDS [Zheng *et al.*, 2009]. However, these datasets are moderately

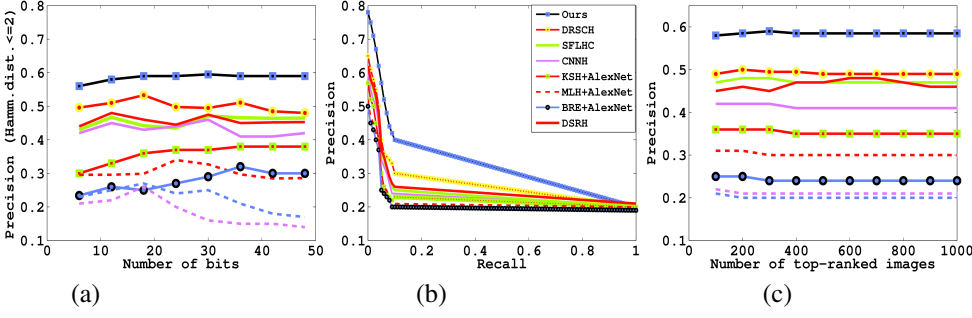


Figure 4: The results on CUHK03. (a) precision curves within Hamming radius; (b) precision-recall curves of Hamming ranking with 48 bits; (c) precision curves with 48 bits with respect to varied number of top-ranked images. Best view in color.




Figure 5: The results on Market-1501. (a) precision curves within Hamming radius; (b) precision-recall curves of Hamming ranking with 48 bits; (c) precision curves with 48 bits with respect to varied number of top-ranked images. Best view in color.

small/medium-sized, rendering them not suitable to be the test bed for our fast hashing learning framework. More recently, to facilitate deep learning in person re-identification, two large datasets *i.e.*, CUHK03 and Market1501 are contributed with more identities captured by multiple cameras in more realistic conditions.

- The **CUHK03** dataset [Li *et al.*, 2014] includes 13,164 images of 1360 pedestrians. The whole dataset is captured with six surveillance camera. Each identity is observed by two disjoint camera views, yielding an average 4.8 images in each view. This dataset provides both manually labeled pedestrian bounding boxes and bounding boxes automatically obtained by running a pedestrian detector [Felzenszwalb *et al.*, 2010]. In our experiment, we report results on labeled data set.
- The **Market-1501** dataset [Zheng *et al.*, 2015] contains 32,643 fully annotated boxes of 1501 pedestrians, making it the largest person re-id dataset to date. Each identity is captured by at most six cameras and boxes of person are obtained by running a state-of-the-art detector, the Deformable Part Model (DPM) [Huang *et al.*, 2015]. The dataset is randomly divided into training and testing sets, containing 750 and 751 identities, respectively.

Competitors We present quantitative evaluations in terms of searching accuracies and compare our method with seven state-of-the-art methods:

- Kernel-based Supervised Hashing (KSH) [Liu *et al.*, 2012]: KSH is a kernel based method that maps the data to binary hash codes by maximizing the separability of code inner products between similar and dissimilar pairs. In particular, KSH adopts the kernel trick to learn nonlinear hash functions on the feature space.
- Minimal Loss Hashing (MLH) [Norouzi and Blei, 2011]: MLS is working by treating the hash codes as latent variables, and employs the structured prediction formulation for hash learning.
- Binary Reconstructive Embedding (BRE) [Kulis and Darrell, 2009]: Without requiring any assumptions on data distributions, BRE directly learns the hash functions by minimizing the reconstruction error between the distances in the original feature space and the Hamming distances in the embedded binary space.
- CNNH [Xia *et al.*, 2014]: is a supervised hashing method in which the learning process is decomposed into a stage of learning approximate hash codes, followed by a second stage of learning hashing functions and image representations from approximate ones.
- Simultaneous Feature Learning and Hash Coding based

Table 2: MAP of Hamming ranking w.r.t. varied number of bits on two person re-identification datasets.

Method	CUHK03				Market-1501			
	24 bits	32 bits	48 bits	128 bits	24 bits	32 bits	48 bits	128 bits
Ours	0.579	0.594	0.602	0.601	0.452	0.466	0.481	0.482
SFLHC [Lai <i>et al.</i> , 2015]	0.428	0.468	0.472	0.476	0.365	0.372	0.377	0.378
DSRH [Zhao <i>et al.</i> , 2015]	0.494	0.486	0.482	0.484	0.403	0.421	0.426	0.423
DRSCH [Zhang <i>et al.</i> , 2015]	0.509	0.517	0.523	0.521	0.421	0.443	0.452	0.447
CNNH [Xia <i>et al.</i> , 2014]	0.403	0.411	0.417	0.414	0.305	0.337	0.346	0.345
KSH+AlexNet [Liu <i>et al.</i> , 2012]	0.301	0.339	0.356	0.357	0.264	0.282	0.288	0.281
MLH+AlexNet [Norouzi and Blei, 2011]	0.262	0.295	0.299	0.302	0.224	0.257	0.269	0.273
BRE+AlexNet [Kulis and Darrell, 2009]	0.206	0.215	0.237	0.239	0.185	0.196	0.211	0.210



Figure 6: The difference between Average Precision (AP) and CMC measurements. The green and red boxes represent the position of true and false matches in rank lists. For all three rank lists, CMC curve at rank 1 remains 1 whilst AP=1 (rank list A), 1 (rank list B), and 0.7 (list C), respectively.

on CNNs (SFLHC) [Lai *et al.*, 2015]: SFLHC is a deep architecture consisting of stacked convolution layers and hashing code learning module. It adopts a triplet ranking loss to preserve relative similarities.

- Deep Semantic Ranking Hashing (DSRH) [Zhao *et al.*, 2015]: DSRH is a recently developed method that incorporates deep feature learning into hash framework in order to preserve multi-level semantic similarity between multi-label images. Also, their network is optimized on a triplet ranking embedding.
- Deep Regularized Similarity Comparison Hashing (DRSCH) [Zhang *et al.*, 2015]: DRSCH is a deep framework which aims to generate bit-scalabel hash codes directly from raw images. Their network is optimized by triplet ranking loss, and hash codes are regularized by adjacency consistency.

The first three methods are conventional supervised methods and the last three are based on deep learning framework. The results of these competitors are obtained by the implementations provided by their authors. For fair comparison on three supervised methods *i.e.*, KSH, MLH, and BRE, we extract CNN features for person images using AlexNet [Krizhevsky *et al.*, 2012], and feed the feature vectors from the last fully-connected layer (4096-dim) to MLH and BRE, denoted as KSH+AlexNet, MLH+AlexNet, BRE+AlexNet, respectively.

Evaluation Protocol We adopt four evaluation metrics in the experiments: Mean Average Precision (MAP), Precision curves with Hamming distance within 2, Precision-Recall curves, and Precision curves with respect to varied number of top returned samples.

In person re-identification, a standard evaluation metric is Cumulated Matching Characteristics (CMC) curve, which shows the probability that a correct match to the query identity appears in different-sized candidate lists. This measurement is, however, is valid only in the single-shot setting where there is only one ground truth match for a given query (see an example in Fig.6). In the case of one-shot, precision and recall are degenerated to be the same manner. Nonetheless, given multiple ground truths regarding to a query identity, the CMC curve is biased due to the fact that the recall issue is not considered. For instance, two rank lists A and B in Fig.6 can yield their CMC value equal to 1 at rank=1, respectively, whereas CMC curves fail to provide a fair comparison of the quality between the two rank lists. By contrast, Average Precision (AP) can quantitatively evaluate the quality of rank list for the case of multi-ground-truth.

For Market-1501 (CUHK03) dataset, there are on average 14.8 (4.8) cross-camera ground truths for each query. Thus, we employ Mean Average Precision (MAP) to evaluate the overall performance. For each query, we calculate the area under the Precision-Recall curve, which is known as Average Precision (AP). Then, MAP is calculated as the mean value of APs over all queries. We have the definition of MAP in the following

$$MAP(Q) = \frac{1}{Q} \sum_{j=1}^{|Q|} \frac{1}{m_j} \sum_{k=1}^{m_j} Precision(R_{jk}), \quad (11)$$

where Q denotes a set of queries, and $\{d_1, \dots, d_{m_j}\}$ are a set of relevant items with respect to a given query $q_j \in Q$. R_{jk} is the set of ranked retrieval results from the top results until item d_k is retrieved.

Given a query, the precision with hamming distance within 2 (@ r -bits) w.r.t. the returned top N nearest neighbors is defined as

$$Precision(\|\cdot\|_H \leq 2)@N = \frac{\#(imgs \cap \{imgs - query\|_H \leq 2\})}{N} \quad (12)$$

where $imgs$ denote similar images to the query, the hamming distance between two binary vectors is the number of

coefficients where they differ. The four types of metrics are widely used to evaluate hashing models [Liu *et al.*, 2012; Lai *et al.*, 2015].

Implementation Details We implemented our architecture using the Theano [Bergstra *et al.*, 2010] deep learning framework with contrastive, triplet, and the proposed structured loss. The batch size is set to 128 for contrastive and our method and to 120 for triplet. Network training converges in roughly 22-24 hours on NVIDIA GTX980. All training and test images are normalized to 160 by 60. We augment the training data by performing random 2D translation, as also done in [Li *et al.*, 2014; Ahmed *et al.*, 2015]. For an original image of size $W \times H$, we sample 5 images around the image center, with translation drawn from a uniform distribution in the range $[-0.05H, 0.05H] \times [-0.05W, 0.05W]$. In training, we exhaustively use all the positive pairs of examples and randomly generate approximately equal number of negative pairs as positives.

In Market-1501, there are 12,936 images for training and 19,732 images for test, corresponding to 750 and 751 identities, respectively. In CUHK03 dataset, we randomly partition the dataset into training, validation, and test with 1160, 100, and 100 identities, respectively. During testing, for each identity, we select one query image in each camera. The search process is performed in a *cross-camera* mode, that is, relevant images captured in the same camera as the query are regarded as “junk” [Philbin *et al.*, 2007], which means that this image has no influence to re-identification accuracy. In this scenario, for Market-1501 dataset, each identity has at most 6 queries, and there are 3,363 query images in total. For CUHK03 dataset, each identity has at most 2 queries, and there are 200 query images in total.

In our implementation, we use all positive anchor positive pairs regarding to each identity. In pairwise training, anchor negative pairs are generated by randomly selecting a sample from a different identity with respect to the anchor’s identity. The same sampling scheme is applied on triplet selection. To add meaningful hard negatives into mini-batch in our model, we select hard neighbors from gallery view for each training image in a positive pair. Specifically, for an anchor \mathcal{I} and its positive \mathcal{I}^+ , their hard negatives \mathcal{I}^- s are selected such that $\|s_{\mathcal{I}} - s_{\mathcal{I}^+}\|_2^2 < \|s_{\mathcal{I}} - s_{\mathcal{I}^-}\|_2^2$, where $s_{(\cdot)}$ is a visual descriptor and in our experiment we use SIFT features at the beginning of training³. Since features are updated continuously as network is on training, $s_{(\cdot)}$ corresponds to feature extracted after each 50 epochs.

4.2 Results on Benchmark Datasets

We test and compare the search accuracies of all methods against two datasets. Comparison results are reported in Table 2 and Figs. 4–5. We can see that

³To extract SIFT/LAB features, we first rescale pedestrian images to a resolution of 48×128 in order to remove some background regions. SIFT and color features are extracted over a set of 14 dense overlapping 32×32 -pixels regions with a step stride of 16 pixels in both directions. Thus, we have the feature descriptor vector with length 9408.

Table 3: Comparison results of skip layer (FC1+FC2) against second fully connected layer (FC2) on two datasets of CUHK03 and Market-1501.

Method	24 bits	32 bits	48 bits	128 bits
MAP (CUHK03)				
FC2	0.529	0.546	0.571	0.584
FC1+FC2	0.579	0.594	0.602	0.601
MAP (Market-1501)				
FC2	0.417	0.420	0.439	0.437
FC1+FC2	0.452	0.466	0.481	0.482

- On the two benchmark datasets, the proposed method outperforms all supervised learning baselines using CNN features in terms of MAP, precision with Hamming distance 2, precision-recall, and precision with varying size of top returned images. For instance, compared with KSH + AlexNet, the MAP results of the proposed method achieves a gain from 35.6% 58.5%, 28.8% 48.1% with 48 bits on CUHK03 and Market-1501, respectively.
- Comparing with CNNH [Xia *et al.*, 2014], which is a two-stage deep network based hashing method, our method indicates a better searching accuracies. Specifically, the MAP results achieve a relative increase by a margin of 16% and 13% on two datasets, respectively. This observation can verify that jointly learning features and hashing codes are beneficial to each other.
- Comparing with the most related competitors DSRH [Zhao *et al.*, 2015] and DRSC [Zhang *et al.*, 2015], our structured prediction suits well to SGD and thus achieves superior performance. For example, in terms of MAP on CUHK03 dataset, a notable improvement can be seen from 49.4% (50.9%) to 54.7%, compared with DSRH [Zhao *et al.*, 2015] (DRSC [Zhang *et al.*, 2015]).

We also conduct self-evaluation of our architecture with skip layer connected to hash layers and its alternative with only the second fully connected layer. As can be seen in Table 3, the results of the proposed architecture outperforms its alternative with only the second fully connected layer as input to the hash layer. One possible reason is the hash layer can see multi-scale features by connecting to the first and second fully connected layers (features in the FC2 is more global than those in FC1). And adding this bypass connections can reduce the possible information loss in the network.

4.3 Comparison with State-of-the-art Approaches

In this section, we evaluate our method by comparing with state-of-the-art approaches in person re-identification. Apart from the above hashing methods, seven competitors are included in our experiment, which are FPNN [Li *et al.*, 2014], JointRe-id [Ahmed *et al.*, 2015], KISSME [Kostinger *et al.*, 2012], SDALF [Farenzena *et al.*, 2010], eSDC [Zhao *et al.*, 2013], kLFDA [Xiong *et al.*, 2014], XQDA [Liao *et al.*, 2015], DomainDropout [Xiao *et al.*, 2016], NullSpace [Zhang *et al.*, 2016] and BoW [Zheng *et al.*, 2015]. For KISSME [Kostinger *et al.*, 2012], SDALF [Farenzena *et al.*,

2010], eSDC [Zhao *et al.*, 2013], kLFDA [Xiong *et al.*, 2014] and BoW model [Zheng *et al.*, 2015], the experimental results are generated by their suggested features and parameter settings. For XQDA [Liao *et al.*, 2015] and NullSpace [Zhang *et al.*, 2016], the Local Maximal Occurrence (LOMO) features are used for person representation. The descriptor has 26,960 dimensions. FPN [Li *et al.*, 2014] is a deep learning method with the validation set adopted to select parameters of the network. JointRe-id [Ahmed *et al.*, 2015] is an improved deep learning architecture in an attempt to simultaneously learn features and a corresponding similarity metric for person re-identification. DomainDropout [Xiao *et al.*, 2016] presents a framework for learning deep feature representations from multiple domains with CNNs. We also extract the intermediate features from the last fully-connected layer, denoted as Ours (FC), to evaluate the performance without hash layer. To have fair comparison with DomainDropout [Xiao *et al.*, 2016], we particularly leverage training data from CUHK03, CUHK01 [Li *et al.*, 2012] with domain-aware dropout, and Market-1501, denoted as Ours (DomainDropout).

Table 4 displays comparison results with state-of-the-art approaches, where all of the Cumulative Matching Characteristics (CMC) Curves are single-shot results on CUHK03 dataset whilst multiple-shot on Market1501 dataset. All hashing methods perform using 128 bits hashing codes, and the ranking list is based on the Hamming distance. We can see that on Market-1501 dataset our method outperforms all baselines on rank 1 recognition rate except NullSpace [Zhang *et al.*, 2016]. The superiority of NullSpace [Zhang *et al.*, 2016] on Market-1501 comes from enough samples in each identity, which allows it to learn a discriminative subspace. Our result (48.06%) is very comparative to NullSpace [Zhang *et al.*, 2016] (55.43%) while the time cost is tremendously reduced, as shown in Table 5. Besides, the performance of our model without hash layer (Ours (FC)) is consistently better than that with hashing projection. This is mainly because the dimension reduction in hashing layer and quantization bring about certain information loss.

On CUHK03 dataset, DomainDropout [Xiao *et al.*, 2016] achieves the best performance in recognition rate at rank from 1 to 10. This is mainly because DomainDropout [Xiao *et al.*, 2016] introduces a method to jointly utilize all datasets in person re-identification to produce generic feature representation. However, this action renders their model extremely expensive in training given a variety of datasets varied in size and distributions. To this end, we test the average testing time of our model and competing deep learning methods, and report results in Table 5. The testing time aggregates computational cost in feature extraction, hash code generation, and image search. For all the experiments, we assume that every image in the database has been represented by its binary hash codes. In this manner, the time consumption of feature extraction and hash code generation is mainly caused by the query image. It is obvious that our model achieves comparable performance in terms of efficiency in matching pedestrian images. Our framework runs slightly slower than DRSCH and SFLHC due to the computation of structured loss on each mini-batch.

4.4 Convergence Study

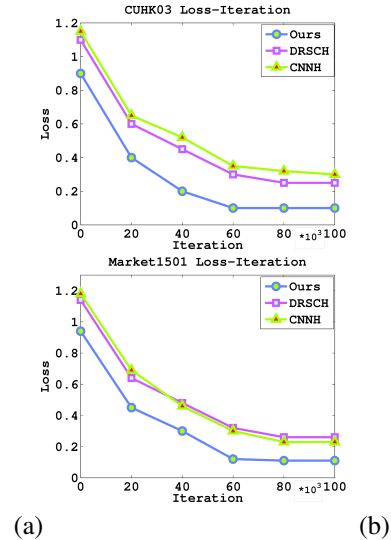


Figure 7: Convergence study on two benchmark datasets. It is obvious that our structured embedding has fast convergence compared with contrastive and triplet embeddings.

In this experiment, we study the convergence speed of optimizing contrastive, triplet, and structured embedding, respectively. The average loss values over all mini-batches are computed on three kinds of embeddings, as shown in Fig.7. We can see that the proposed structured embedding is able to converge faster than other two embeddings. This can be regarded as the response to the augment from hard negatives which provide informative gradient direction for positives.

5 Conclusion

In this paper, we developed a structured deep hashing architecture for efficient person re-identification, which jointly learn both CNN features and hash functions/codes. As a result, person re-identification can be resolved by efficiently computing and ranking the Hamming distances between images. A structured loss function is proposed to achieve fast convergence and more stable optimization solutions. Empirical studies on two larger benchmark data sets demonstrate the efficacy of our method. In our future work, we would explore more efficient training strategies to reduce training complexity, and possible solutions include an improved loss function based on local distributions.

References

- [Ahmed *et al.*, 2015] Ejaz Ahmed, Michael Jones, and Tim K. Marks. An improved deep learning architecture for person re-identification. In *CVPR*, 2015.
- [Bergstra *et al.*, 2010] James Bergstra, Olivier Breuleux, Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, Guillaume Desjardins, Joseph Turian, David Warde-Farley, and Yoshua Bengio. Theano: a cpu and gpu math expression compiler. In *Proceedings of the Python for Scientific Computing Conference (SciPy)*, 2010.

Table 4: Comparison with state-of-the-art approaches on two person re-identification datasets. The evaluation is based on CMC. “-” indicates that the result is not available.

Method	CUHK03 (CMC %)				Market-1501 (CMC %)			
	r=1	r=5	r=10	r=20	r=1	r=5	r=10	r=20
Ours	37.41	61.28	77.46	88.42	48.06	61.23	75.67	87.06
Ours (FC2)	43.20	69.07	84.24	90.92	50.12	63.50	76.82	89.24
Ours (DomainDropout)	74.20	92.27	94.24	95.92	58.12	68.50	80.82	92.24
SFLHC [Lai <i>et al.</i> , 2015]	12.38	30.52	49.34	71.55	37.74	59.09	74.25	86.52
DSRH [Zhao <i>et al.</i> , 2015]	9.75	28.10	47.82	67.95	34.33	59.82	71.27	86.09
DRSCH [Zhang <i>et al.</i> , 2015]	20.84	49.39	72.66	83.03	41.25	58.98	76.04	85.33
CNNH [Xia <i>et al.</i> , 2014]	8.27	22.53	45.09	59.74	16.46	39.95	51.24	71.23
KSH+AlexNet [Liu <i>et al.</i> , 2012]	5.65	15.71	22.75	34.68	12.57	31.22	48.66	66.72
MLH+AlexNet [Norouzi and Blei, 2011]	5.75	15.62	27.61	42.68	10.89	29.93	46.78	66.32
BRE+AlexNet [Kulis and Darrell, 2009]	4.91	11.24	17.83	29.20	12.65	32.68	49.08	67.70
FPNN [Li <i>et al.</i> , 2014]	20.65	50.09	66.42	80.02	-	-	-	-
JointRe-id [Ahmed <i>et al.</i> , 2015]	54.74	86.71	91.10	97.21	-	-	-	-
KISSME [Kostinger <i>et al.</i> , 2012]	14.17	41.12	54.89	70.09	40.47	59.35	75.26	83.58
SDALF [Farenzena <i>et al.</i> , 2010]	5.60	23.45	36.09	51.96	20.53	47.82	62.45	80.21
eSDC [Zhao <i>et al.</i> , 2013]	8.76	27.03	38.32	55.06	33.54	58.25	74.33	84.57
kLFDA [Xiong <i>et al.</i> , 2014]	47.25	64.58	82.36	89.17	-	-	-	-
BoW [Zheng <i>et al.</i> , 2015]	24.33	58.42	71.28	84.91	47.25	62.71	75.33	86.42
DomainDropout [Xiao <i>et al.</i> , 2016]	72.60	92.22	94.50	95.01	-	-	-	-
NullSpace (LOMO [Liao <i>et al.</i> , 2015]) [Zhang <i>et al.</i> , 2016]	58.90	85.60	92.45	96.30	55.43	-	-	-
XQDA (LOMO) [Liao <i>et al.</i> , 2015]	52.20	82.23	92.14	96.25	43.79	-	-	-

Table 5: Comparison on the average testing time (millisecond per image) by fixing the code length to be 48 (128) bits. ∞ indicates that the computational cost is too high to be estimated and “-” indicates that the result is not available.

Method	CUHK03 (ms)	Market-1501 (ms)
Ours	4.617 (4.982)	7.374 (7.902)
SFLHC [Lai <i>et al.</i> , 2015]	4.241 (4.782)	5.892 (6.417)
DSRH [Zhao <i>et al.</i> , 2015]	5.765 (6.019)	7.887 (8.445)
DRSCH [Zhang <i>et al.</i> , 2015]	2.332 (2.816)	3.609 (3.973)
CNNH [Xia <i>et al.</i> , 2014]	5.359 (5.743)	6.943 (7.410)
KSH+AlexNet [Liu <i>et al.</i> , 2012]	7.279 (7.805)	9.046 (9.537)
MLH+AlexNet [Norouzi and Blei, 2011]	6.727 (7.198)	8.092 (8.545)
BRE+AlexNet [Kulis and Darrell, 2009]	6.765 (7.214)	9.072 (9.613)
FPNN [Li <i>et al.</i> , 2014]	∞	∞
JointRe-id [Ahmed <i>et al.</i> , 2015]	∞	∞
DomainDropout [Xiao <i>et al.</i> , 2016]	∞	∞
kLFDA [Xiong <i>et al.</i> , 2014]	-	43.4×10^3
XQDA [Liao <i>et al.</i> , 2015]	-	1.6×10^3
MFA [Xiong <i>et al.</i> , 2014]	-	43.2×10^3
NullSpace [Zhang <i>et al.</i> , 2016]	-	31.3×10^3

[Chen *et al.*, 2016] Shi-Zhe Chen, Chun-Chao Guo, and Jian-Huang Lai. Deep ranking for re-identification via joint representation learning. *IEEE Transactions on Image Processing*, 25(5):2353–2367, 2016.

[Ding *et al.*, 2015] Shengyong Ding, Liang Lin, Guangrun Wang, and Hongyang Chao. Deep feature learning with relative distance comparison for person re-identification. *Pattern Recognition*, 48:2993–3003, 2015.

[Farenzena *et al.*, 2010] M. Farenzena, L. Bazzani, A. Perina, V. Murino, and M. Cristani. Person re-identification by symmetry-driven accumulation of local features. In *CVPR*, 2010.

[Felzenszwalb *et al.*, 2010] P.F. Felzenszwalb, R.B. Girshick, D. McAllester, and D. Ramanan. Object detec-

tion with discriminatively trained part-based models. *IEEE TPAMI*, 32(9):1627–1645, 2010.

[Gong and Lazebnik, 2011] Y. Gong and S. Lazebnik. Iterative quantization: A procrustean approach to learning binary codes. In *CVPR*, pages 817–824, 2011.

[Gong *et al.*, 2012] Y. Gong, S. Kumar, V. Verma, and S. Lazebnik. Angular quantization-based binary codes for fast similarity search. In *NIPS*, pages 1196–1204, 2012.

[Gong *et al.*, 2014] Shaogang Gong, Marco Christani, Chen Changee Loy, and Timothy M. Hospedales. *Person Re-Identification*. Springer, 2014.

[Gray and Tao, 2008] D. Gray and H. Tao. Viewpoint invariant pedestrian recognition with an ensemble of localized features. In *ECCV*, 2008.

[Gray *et al.*, 2007] D. Gray, S. Brennan, and H. Tao. Evaluating appearance models for recognition, reacquisition, and tracking. In *Proc. Int’l. Workshop on Perf. Eval. of Track. and Surv’l.*, 2007.

[He *et al.*, 2010] J. He, W. Liu, and S.F. Chang. Scalabel similarity search with optimized kernel hashing. In *KDD*, pages 1129–1138, 2010.

[He *et al.*, 2013] K. He, F. Wen, and J. Sun. K-means hashing: an affinity-perserving quantization method for learning binary compact codes. In *CVPR*, pages 2938–2945, 2013.

[Hirzer *et al.*, 2011] M. Hirzer, C. Belezna, P. M. Roth, and H. Bischof. Person re-identification by descriptive and discriminative classification. In *Proc. Scandinavian Conf. on Image Anal.*, 2011.

- [Hoffer and Ailon, 2014] Elad Hoffer and Nir Ailon. Deep metric learning using triplet network. In *arXiv preprint arXiv: 1412.6622*, 2014.
- [Huang *et al.*, 2015] B. Huang, J. Chen, Y. Wang, C. Liang, Z. Wang, and K. Sun. Sparsity-based occlusion handling method for person re-identification. In *Multimedia Modeling*, 2015.
- [Joachims *et al.*, 2009] Thorsten Joachims, Thomas Finley, and Chun-Nam John Yu. Cutting-plane training of structural svms. *Journal of machine learning research*, 77:27–59, 2009.
- [Kostinger *et al.*, 2012] M. Kostinger, M. Hirzer, P. Wohlhart, P. M. Roth, and H. Bischof. Large scale metric learning from equivalence constraints. In *CVPR*, 2012.
- [Krizhevsky *et al.*, 2012] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [Kulis and Darrell, 2009] B. Kulis and T. Darrell. Learning to hash with binary reconstructive embeddings. In *NIPS*, pages 1042–1050, 2009.
- [Lai *et al.*, 2015] Hanjiang Lai, Yan Pan, Ye Liu, and Shuicheng Yan. Simultaneous feature learning and hash coding with deep neural networks. In *CVPR*, 2015.
- [Li *et al.*, 2012] Wei Li, Rui Zhao, and Xiaogang Wang. Human identification with transferred metric learning. In *ACCV*, 2012.
- [Li *et al.*, 2013] Z. Li, S. Chang, F. Liang, T. S. Huang, L. Cao, and J.R. Smith. Learning locally-adaptive decision functions for person verification. In *CVPR*, 2013.
- [Li *et al.*, 2014] W. Li, R. Zhao, T. Xiao, and X. Wang. Deepreid: Deep filter pairing neural network for person re-identification. In *CVPR*, 2014.
- [Liao *et al.*, 2015] Shengcai Liao, Yang Hu, Xiangyu Zhu, and Stan Z. Li. Person re-identification by local maximal occurrence representation and metric learning. In *CVPR*, 2015.
- [Liong *et al.*, 2015] Venice Erin Liong, Jiwen Lu, Gang Wang, Pierre Moulin, and Jie Zhou. Deep hashing for compact binary codes learning. In *CVPR*, pages 2475–2483, June 2015.
- [Liu *et al.*, 2012] Wei Liu, Jun Wang, Rongrong Ji, Yu-Gang Jiang, and Shih-Fu Chang. Supervised hashing with kernels. In *CVPR*, pages 2074–2081, 2012.
- [Norouzi and Blei, 2011] M. Norouzi and D. M. Blei. Minimal loss hash for compact binary codes. In *ICML*, pages 353–360, 2011.
- [Paisitkriangkrai *et al.*, 2015] Sakrapee Paisitkriangkrai, Chunhua Shen, and Anton van den Hengel. Learning to rank in person re-identification with metric ensembles. In *CVPR*, 2015.
- [Pedagadi *et al.*, 2013] S. Pedagadi, J. Orwell, S. Velastin, and B. Boghossian. Local fisher discriminant analysis for pedestrian re-identification. In *CVPR*, 2013.
- [Philbin *et al.*, 2007] James Philbin, Ondrej Chum, Michael Isard, Josef Sivic, and Andrew Zisserman. Object retrieval with large vocabularies and fast spatial matching. In *CVPR*, pages 1–8, 2007.
- [Song *et al.*, 2016] Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. Deep metric learning via lifted structured feature embedding. In *CVPR*, 2016.
- [Sun *et al.*, 2014] Yi Sun, Xiaogang Wang, and Xiaoou Tang. Deep learning face representation from predicting 10000 classes. In *CVPR*, 2014.
- [Sunderrajan and Manjunath, 2016] Santhoshkumar Sunderrajan and B. S. Manjunath. Context-aware hypergraph modeling for person re-identification and summarization. *IEEE Transactions on Multimedia*, 18(1):51–63, 2016.
- [Tsochantaridis *et al.*, 2004] Ioannis Tsochantaridis, Thomas Hofman, Thorsten Joachims, and Yasemin Altun. Support vector machine learning for interdependent and structured output spaces. In *ICML*, 2004.
- [Wang *et al.*, 2007] X. Wang, G. Doretto, T. Sebastian, J. Rittscher, and P. Tu. Shape and appearance context modeling. In *ICCV*, 2007.
- [Wang *et al.*, 2013] Yang Wang, Xuemin Lin, and Qing Zhang. Towards metric fusion on multi-view data: a cross-view based graph random walk approach. In *ACM CIKM*, pages 805–810, 2013.
- [Wang *et al.*, 2015a] Yang Wang, Xuemin Lin, Lin Wu, and Wenjie Zhang. Effective multi-query expansions: Robust landmark retrieval. In *ACM Multimedia*, pages 79–88, 2015.
- [Wang *et al.*, 2015b] Yang Wang, Xuemin Lin, Lin Wu, Wenjie Zhang, and Qing Zhang. Lbmch: Learning bridging mapping for cross-modal hashing. In *ACM SIGIR*, 2015.
- [Wang *et al.*, 2015c] Yang Wang, Xuemin Lin, Lin Wu, Wenjie Zhang, Qing Zhang, and Xiaodi Huang. Robust subspace clustering for multi-view data by exploiting correlation consensus. *IEEE Trans. Image Processing*, 24(11):3939–3949, 2015.
- [Wang *et al.*, 2016a] Yang Wang, Xuemin Lin, Lin Wu, Qing Zhang, and Wenjie Zhang. Shifting multi-hypergraphs via collaborative probabilistic voting. *Knowl. Inf. Syst.*, 46(3):515–536, 2016.
- [Wang *et al.*, 2016b] Yang Wang, Wenjie Zhang, Lin Wu, Xuemin Lin, Meng Fang, and Shirui Pan. Iterative views agreement: An iterative low-rank based structured optimization method to multi-view spectral clustering. In *IJ-CAI*, pages 2153–2159, 2016.
- [Wang *et al.*, 2016c] Zheng Wang, Ruimin Hu, Chao Liang, Yi Yu, Junjun Jiang, Meng Ye, Jun Chen, and Qingming Leng. Zero-shot re-identification via cross-view consistency. *IEEE Transactions on Multimedia*, 18(2):260–272, 2016.
- [Wang *et al.*, 2017a] Yang Wang, Xuemin Lin, Lin Wu, and Wenjie Zhang. Effective multi-query expansions:

- Collaborative deep networks for robust landmark retrieval. *IEEE Transactions on Image Processing*, doi: 10.1109/TIP.2017.2655449, 2017.
- [Wang *et al.*, 2017b] Yang Wang, Wenjie Zhang, Lin Wu, Xuemin Lin, and Xiang Zhao. Unsupervised metric fusion over multiview data by graph random walk-based cross-view diffusion. *IEEE Trans. Neural Netw. Learning Syst.*, 28(1):57–70, 2017.
- [Weiss *et al.*, 2008] Y. Weiss, A. Torralba, and R. Fergus. Spectral hashing. In *NIPS*, 2008.
- [Wilson and Martinez, 2003] D.R. Wilson and T.R. Martinez. The general inefficiency of batch training for gradient decent learning. *Neural Networks*, 16(10):1429–1451, 2003.
- [Wu and Wang, 2017] Lin Wu and Yang Wang. Robust hashing for multi-view data: Jointly learning low-rank kernelized similarity consensus and hash functions. *Image and Vision Computing*, 57:58–66, 2017.
- [Wu *et al.*, 2013] Lin Wu, Yang Wang, and John Shepherd. Efficient image and tag co-ranking: a bregman divergence optimization method. In *ACM Multimedia*, 2013.
- [Wu *et al.*, 2016a] Lin Wu, Chunhua Shen, and Anton van den Hengel. Deep linear discriminant analysis on fisher networks: A hybrid architecture for person re-identification. *Pattern Recognition*, 2016.
- [Wu *et al.*, 2016b] Lin Wu, Chunhua Shen, and Anton van den Hengel. Personnet: Person re-identification with deep convolutional neural networks. In *CoRR abs/1601.07255*, 2016.
- [Xia *et al.*, 2014] R. Xia, Y. Pan, H. Lai, C. Liu, and S. Yan. Supervised hashing for image retrieval via image representation learning. In *AAAI*, 2014.
- [Xiao *et al.*, 2016] Tong Xiao, Hongsheng Li, Wanli Ouyang, and Xiaogang Wang. Learning deep feature representation with domain guided dropout for person re-identification. In *CVPR*, 2016.
- [Xiong *et al.*, 2014] F. Xiong, M. Gou, O. Camps, and M. Szaier. Person re-identification using kernel-based metric learning methods. In *ECCV*, 2014.
- [Yang *et al.*, 2014] Y. Yang, J. Yang, J. Yan, S. Liao, D. Yi, and S.Z. Li. Salient color names for person re-identification. In *ECCV*, pages 536–551, 2014.
- [Ye *et al.*, 2016] Mang Ye, Chao Liang, Yi Yu, Zheng Wang, Qingming Leng, Chunxia Xiao, Jun Chen, and Ruimin Hu. Person re-identification via ranking aggregation of similarity pulling and dissimilarity pushing. *IEEE Transactions on Multimedia*, (99):1–1, 2016.
- [Yi *et al.*, 2014] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z. Li. Deep metric learning for person re-identification. In *ICPR*, 2014.
- [Zhang *et al.*, 2015] Ruimao Zhang, Liang Lin, Wangmeng Zuo, and Lei Zhang. Bit-scalable deep hashing with regularized similarity learning for image retrieval and person re-identification. *IEEE Transactions on Image Processing*, 24(12):4766–4779, December 2015.
- [Zhang *et al.*, 2016] Li Zhang, Tao Xiang, and Shaogang Gong. Learning a discriminative null space for person re-identification. In *CVPR*, 2016.
- [Zhao *et al.*, 2013] Rui Zhao, Wanli Ouyang, and Xiaogang Wang. Unsupervised salience learning for person re-identification. In *CVPR*, 2013.
- [Zhao *et al.*, 2014] Rui Zhao, Wanli Ouyang, and Xiaogang Wang. Learning mid-level filters for person re-identification. In *CVPR*, pages 2528–2535, 2014.
- [Zhao *et al.*, 2015] Fang Zhao, Yongzhen Huang, Liang Wang, and Tieniu Tan. Deep semantic ranking based hashing for multi-label image retrieval. In *CVPR*, 2015.
- [Zheng *et al.*, 2009] Wei-Shi Zheng, Shaogang Gong, and Tao Xiang. Associating groups of people. In *BMVC*, 2009.
- [Zheng *et al.*, 2015] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *ICCV*, 2015.